

Modeling The Term Structure of Commodity Futures Curve and Empirical Study

David Lee

Abstract

Commodity derivatives trading typically depends on futures/forwards. As such, there is a need to model the dynamics of futures curves. This article presents a model for valuing commodity derivatives by decomposing the futures term structure into its principal components and studying their economic drivers. We also propose generic online and offline approaches for model calibration. Empirical study shows that the model can better capture the commodity seasonality, price dynamics, and volatility skew/smile.

Key words: commodity derivatives, multiple factor model, model calibration, volatility skew, volatility smile, empirical study.

JEL Classification: E44, G21, G12, G24, G32, G33, G18, G28

1. Introduction

Commodity futures/forward curves are often viewed as the market expectation for the evolution of spot prices in the future. Understanding the behavior of commodity futures curves is vital for commodity trading and hedging. First it provides valuable information about the market trends and expectations. Second it helps to relate future expected prices with today's actual prices. Third, it may shed light on some economic policy implications.

Some commodities, such as natural gas, exhibit seasonality in prices. For example, in the winter months, the higher demand for natural gas raises the prices. Therefore, commodity prices are simulated with a seasonal average and a random component.

Commodity futures are often modeled without a drift. This is because a very large amount of historical data would be required to calibrate any sort of reasonable drift. In general, commodities rise and fall over time, and it is generally difficult to associate a drift to such motion.

The random component of the commodity future prices is generally broken down into major contributors or factors. These are known as the principal components (PCs). The principal components for the same commodity are generally independent – that is, we assume that there exists in the world N major independent factors driving the randomness in the price.

There is a vast literature on commodity modeling. Kirk (1995) finds that the randomness contributed by the principal components on the commodity prices are correlated. The complex stochastic behavior in commodity prices cannot be fully described by one-factor models. Schwartz and Smith (2000) proposed a two-factor model that assumes the first factor to be a zero-mean Ornstein–Uhlenbeck process, representing short term fluctuations of price. The second factor is modeled by the arithmetic Brownian Motion, representing a long-term effect. Both factors are assumed to be correlated.

Borovkova and Geman (2006, 2007) introduce a synthetic factor model that prices options in a robust way. The approach approximates the distribution of a basket option based on the moment matching

technique. Ladokhin & Borovkova (2021) extend Borovkova and Geman (2007) into a three-factor model that incorporates the synthetic spot price. The model is based on liquidly traded futures and stochastic level of mean reversion.

Hu et al. (2020) find that there is a strong correlation between commodity prices and equity markets, and lead to a breakdown of the obvious relationship between commodities and cyclical phases of the economy. Idilbi-Bayaa and Qadan (2021) study the link between yield spread and long-term prices in the commodity market and use yield curve to forecast commodity prices.

Fasanya and Awodimila (2020) indicate that long-term predictions about commodity prices are essential in targeting inflation and promoting overall economic stability. Hollstein et al. (2020) examine the volatility term structures in commodity markets and model state-dependent spillovers in principal components of the volatility term structures of different commodities, as well as that of the equity market.

This article presents a new model for commodity derivative valuation. The model clearly illustrates the essence of the multi-factor driving of commodity prices. Commodity price is represented as a sum of mean reverting factors. Factors with low mean reversion rate are called Slow/Long Factors that represent uncertainties on long-term fundamentals. Factors with high mean reversion rate are called Fast/Short Factors that capture short term risk factors such as weather, Middle East unrest, economic slowdown, central-bank policy changes, etc.

The general effect of each principal component on the price is random and thus in a simulation, a correlated random number is generated for each principal component. These random numbers are known as the Fourier coefficients. The final price for a given contract is then simply its seasonal mean plus the random Fourier coefficient multiplied by its corresponding principal component.

Each of the principal components are column vectors containing as many elements as contracts in the term structure. The magnitude of the elements of the vector corresponding to the earlier contracts, such as one, two, three, etc. month contracts are usually higher than later contracts,

such as the 30-month contract, which increases the volatility of earlier contracts. This correctly models the observation that shorter contracts normally exhibit higher volatility than longer contracts.

We also propose a generic model calibration procedure. The calibration procedure consists of an offline step where the mean reversion rates, the ratio of the long and short factor volatilities, and the correlation between the long and short factors are determined via historical analysis. This offline step is performed relatively infrequently (typically once a month). There's also an online step of the calibration which happens every time the model is used to price an option or to simulate price paths.

The online calibration makes sure that the model prices for vanilla options match exactly the market prices. Finally, there is an introspective step where trading desks adjust the factor parameters in the offline step to better match the market prices for non-vanilla options such as swaptions and term-settled options. This introspective step happens whenever we have an observation of the market price for these illiquid options, through broker quotes or market consensus services.

The rest of this article is organized as follows: First we present the commodity factor model. Second, we describe the generic calibration method for the model. Third, we elaborate the valuation of commodity derivatives and Monto-Carlo simulation. Next, we conduct an empirical study and discuss the result. Finally, the conclusions are provided.

2. The Model

A single factor model assumes that the forward curve of an asset is driven entirely by the spot price. The generic formulation reduces to

$$\frac{dF^T(t)}{F^T(t)} = p(t)dW(t)$$

The function $p(t)$ is the familiar time-dependent local volatility function, which is typically calibrated via bootstrapping to the term structure of market implied volatility for options at different maturities.

Data analysis on historical crude oil futures prices revealed that more than 90% of the variance of the forward curve can be captured by the first two principal components. Hence, it's customary to use a two-factor model to describe the dynamics of crude oil futures curve. Since crude oil is easy to store and globally traded, it is less subject to seasonal fluctuations than natural gas or electricity.

Crude oil futures with different delivery months are intrinsically linked and traditionally people simplify the model by dropping the T-dependent volatility function q . To reduce the number of parameters, we also assume that the mean reversion rate of the long factor is zero, leaving only a single mean reversion rate for the short factor to be determined. The model simplifies to

$$\frac{dF^T(t)}{F^T(t)} = p_S(t)e^{-\beta(T-t)}dW_S(t) + p_L(t)dW_L(t)$$

Natural gas and power are strongly seasonal, and such assets delivered at different months, or even at different hours of the day in the case of electricity, have distinct characteristics and can be thought of as different commodities. To reflect this in the model, we utilize the T-dependent volatility q and drop the t-dependent volatility p . The model reduces to:

$$\frac{dF^T(t)}{F^T(t)} = q_S(T)e^{-\beta(T-t)}dW_S(t) + q_L(T)dW_L(t)$$

For financial derivatives such as futures, it mostly suffices to use a two-factor model. For physical natural gas or power (or financial contracts that closely track physical prices such as gas daily index), sometimes we need to add a “super”-fast factor to model the weather induced shocks. For example, on an extremely hot day, the air-conditioning usage may lead to a demand peak in power and the power price can suddenly jump by many folds. The usual two-factor model will not be able to capture this risk occurring only at the time scale of the delivery day or hour.

A three-factor model may also be useful to value the time spread options, particularly a portfolio of time spread options, such as natural gas storage contracts. In the two-factor model, the two principal modes of motion are the parallel curve shift and the curve rotation. Time spread options are not sensitive to the primary mode of curve shift. They are sensitive to curve rotation and other curve shape changes beyond the first two principal components.

Commodity derivatives typically are traded by futures/forwards and there's a need to model the dynamics of the entire forward curve. The application of such a model includes (semi-)analytical pricing of non-vanilla options, Monte-Carlo simulation, risk management, etc. The Multi-factor model can model a wide range of commodity assets, as well as FX and equities.

The model assumes that the forward curve is driven by N random factors, characterized by different time scales. Mathematically, the over-arching formula for the futures price F^T reads

$$\frac{dF^T(t)}{F^T(t)} = \sum_{i=1}^N p_i(t) q_i(T) e^{-\beta_i(T-t)} dW_i(t)$$

Here, t is the current time, T is the futures expiration time, i is the index of the i 'th factor, N is the total number of factors, p_i is the local time dependent volatility of the i 'th factor, q_i is the expiry time dependent volatility of the i 'th factor, β_i is the mean reversion rate of the i 'th factor, and W_i is a standard Brownian motion that is the random driver for the i 'th factor. The Brown motions are correlated with the correlation matrix with element $\rho_{i,j}$, i.e. $\langle dW_i, dW_j \rangle = \rho_{i,j} dt$.

Here are several examples of the generic formulation above, from which the intuition of the model should come alive. But for now, let's expand certain properties of the generic model a bit more. First, we define the i 'th Mean Reverting Factor Y_i by the following OU process.

$$dY_i(t) = -\beta_i Y_i(t) dt + p_i(t) dW_i(t)$$

We assume that the process starts from 0, i.e. $Y_i(0) = 0$. The solution to this SDE is easily derivable:

$$\begin{aligned} d[e^{\beta_i t} Y_i(t)] &= e^{\beta_i t} p_i(t) dW_i(t) \\ Y_i(t) &= \int_0^t e^{-\beta_i(t-s)} p_i(s) dW_i(s) \end{aligned}$$

The futures price can be expressed in terms of the mean reverting factor as follows:

$$\frac{dF^T(t)}{F^T(t)} = \sum_{i=1}^N q_i(T) e^{-\beta_i T} d[e^{\beta_i t} Y_i(t)]$$

$$F^T(t) = F^T(0)e^{\sum_{i=1}^N q_i(T)e^{-\beta_i(T-t)}Y_i(t) - \frac{1}{2}\langle \log F^T \rangle_0^t}$$

Here, $\langle \log F^T \rangle_0^t$ is the quadratic variation of the log futures price from time 0 to time t. The significance of the formula above is that the futures price can be explicitly put in a simple algebraic formula involving the N mean reverting factors which do not depend on T.

The quadratic variation is a specific case of a very useful quantity: Log Covariance. We define the log covariance of two futures contracts, F^{T_1} and F^{T_2} , integrating from t_1 to t_2 as

$$\langle \log F^{T_1}, \log F^{T_2} \rangle_{t_1}^{t_2} = \sum_{i,j=1}^N q_i(T_1)q_j(T_2) \int_{t_1}^{t_2} p_i(s)p_j(s) e^{-\beta_i(T_1-s)-\beta_j(T_2-s)} \rho_{i,j} ds$$

In particular,

$$\langle \log F^T \rangle_0^t = \sum_{i,j=1}^N q_i(T)q_j(T) \int_0^t p_i(s)p_j(s) e^{-(\beta_i+\beta_j)(T-s)} \rho_{i,j} ds$$

The value of a vanilla option is computed by substituting the volatility $\sqrt{\frac{\langle \log F^T \rangle_0^t}{t}}$ into the Black-Scholes formula.

Substituting T = t into the equation above gives the formula for the spot price of a commodity:

$$S(t) = S(0)e^{\sum_{i=1}^N q_i(t)Y_i(t) + \dots}$$

This equation most clearly illustrates the essence of the multi-factor model. The model describes the (log) commodity price as a sum of mean reverting factors. Factors with low mean reversion rate are called Slow/Long Factors that represent uncertainties on long-term fundamentals. Factors with high mean reversion rate are called Fast/Short Factors that are used to model short term risk factors such as weather, middle-east unrest, economic slowdown, central-bank policy changes, etc.

Finally, if we have two assets that are modeled by the multi-factor model, with potentially different number of factors, we can calculate the Cross-Asset Log Covariance as follows.

$$\begin{aligned}\frac{dF_1^T(t)}{F_1^T(t)} &= \sum_{i=1}^{N_1} p_{i,1}(t) q_{i,1}(T) e^{-\beta_{i,1}(T-t)} dW_{i,1}(t) \\ \frac{dF_2^T(t)}{F_2^T(t)} &= \sum_{j=1}^{N_2} p_{j,2}(t) q_{j,2}(T) e^{-\beta_{j,2}(T-t)} dW_{j,2}(t) \\ \langle \log F_1^{T_1}, \log F_2^{T_2} \rangle_{t_1}^{t_2} &= \sum_{i=1}^{N_1} \sum_{j=1}^{N_2} q_{i,1}(T_1) q_{j,2}(T_2) \int_{t_1}^{t_2} p_{i,1}(s) p_{j,2}(s) e^{-\beta_{i,1}(T_1-s) - \beta_{j,2}(T_2-s)} \rho_{i,j}^{1,2} ds\end{aligned}$$

Here, $\rho_{i,j}^{1,2}$ is the instantaneous correlation between the i 'th factor of the first asset and the j 'th factor of the second asset, i.e.

$$\rho_{i,j}^{1,2} dt = \langle dW_{i,1}(t), dW_{j,2}(t) \rangle$$

The standard multi-factor model specifies a log-normal process for the futures price, and therefore does not incorporate volatility smile directly. The model needs to be adjusted to handle the volatility smile. There are several sophisticated approaches doing this.

The most convenient “moneyness” metric is the so-called **Quick Delta** (QD). QD measures how many standard deviations away the forward price is relative to the strike of the option, in the log-normal space. QD approximates the Black-Scholes (put) delta. But it is more convenient than Black-Scholes delta for the purpose of parameterizing the volatility surface because QD for ATM option is exactly 50% and it is easy to convert QD to a dollar strike or the other way around. Mathematically, QD is defined as

$$QD = N(\log \frac{\text{Strike}}{\text{Forward}} / \sigma \sqrt{T})$$

Here, N is the cumulative normal density function, σ is the ATM volatility, and T is the option expiration. The cumulative normal density function is applied to transform the moneyness to a number between 0 and 1.

3. Model Calibration

The goals of the model calibration are to achieve: i) the market prices of vanilla options need to be perfectly matched by the model; ii) the model prices for non-vanilla options are in good agreement with the market prices; and iii) the implied model dynamics are in good agreement with the characteristics of the historical data series.

In this section, we show how the two-factor model parameters are determined in the offline step of the calibration. The purpose of the offline step is to back out the mean reversion rate (β), the volatility ratio (p_S/p_L for non-seasonal model and q_S/q_L for seasonal model), and the correlation (ρ) between the long and short factors from the historical time series. There are of course many ways of studying historical data and it is clearly beneficial to benchmark the model dynamics against history in multiple ways. We show 3 different methods here.

The first method is Principal Component Analysis (PCA) that is a mathematical procedure that uses an orthogonal transformation to convert a set of observations of possibly correlated variables into a set of values of linearly uncorrelated variables called Principal Components. The number of principal components is less than or equal to the number of original variables.

This transformation is defined in such a way that the first principal component has the largest possible variance or accounts for as much of the variability in the data as possible. Each succeeding component in turn has the highest variance possible under the constraint that it be orthogonal to (i.e., uncorrelated with) the preceding components. Principal components are guaranteed to be independent if the data set is jointly normally distributed.

We apply PCA to the historical daily returns of the futures curve and compare the first two principal components with those implied by the two-factor model. The mean reversion rate, volatility ratio and correlation are chosen to achieve the best fit to the historical PCA in the least squared sense.

We obtain data from FinPricing (see <https://finpricing.com/lib/IrBasisCurve.html>). The following graphs illustrate the quality of fitting for West Texas Intermediate (WTI). The historical period is between July 24, 2018 and July 24, 2021, with more recent data given higher importance. The weighting of the data follows an exponential formula with the decay rate corresponding to a half-life of 125 business days. The two-factor model parameters are 0.35 for mean reversion rate, 1.6 for volatility ratio and -20% for correlation.





FIGURE 1 illustrates the quality of fitting for West Texas Intermediate (WTI).

The second method is to directly compare certain statistical metrics between the historical data and the model. As far as the most liquid non-vanilla options are concerned, the most important statistics are the relative volatilities of different futures and the correlations between these futures. The figures below illustrate the quality of fitting in terms of these two metrics. The historical data and two-factor model parameters are the same as above.




FIGURE 2 illustrates the relative volatilities of different futures for West Texas Intermediate (WTI). More specifically, the graph plots the ratio of the instantaneous volatility of the nearby number in the x-axis versus the instantaneous volatility of the 12th nearby.




FIGURE 3 illustrates the correlation between different futures for West Texas Intermediate (WTI). More specifically, the graph plots the instantaneous correlation of the nearby number in the x-axis and the 12th nearby.

The third method is so-called relative value analysis. Both the first and second methods are purely historical analysis. The strength of the relative value analysis is that it combines information from historical analysis and information from market implied volatilities of the vanilla options.

The key quantity that links the past and the future is the ratio between the realized swap volatility and the realized futures volatility. For example, in calendar 2018, the realized volatility ratio between Cal21 swap and the Feb21 future contract was 0.9. Suppose on Jan 2, 2019, we need to price a Cal22 swaption and the Feb22 vanilla option is priced at 20% implied volatility.

The relative value analysis would suggest that we price the Cal22 swaption using $20\% \times 0.9 = 18\%$ volatility. This gives us a price target for the swaption. We can establish the price targets for a number of tenors, Cal23, Cal24, 2H21, 1H22, etc., and then adjust the mean reversion, volatility ratio and correlations to match as well as possible these target prices. Here's an example of doing historical relative value analysis on Cal21 swap using the historical period of 2018.




FIGURE 4 shows historical realized vol vs. swap vol ratio

Note that the two-factor model parameters we backed out from historical analysis are not used directly to price new deals or manage our risks. Historical analysis is backward looking, while derivatives pricing should reflect our forward-looking view. Therefore, we use the introspective step to complement the historical analysis with the information embedded in the market non-vanilla option prices about how the future will hold.

We use the historical two-factor model parameters to price non-vanilla options such as swaptions or term-settled Asian options, and fine tune the parameters so that the model price is in good agreement with the market prices. These non-vanilla trades typically have wide bid-offer in the market, so we don't need to match the market price exactly.

Once the mean reversion rates, the volatility ratio and the correlations are fixed, the only thing left is to impute the overall level of the factor volatilities. The purpose of the online step of the calibration is

to solve for the factor volatilities such that the model prices vanilla options at exactly their market prices. We now return to the general multi-factor model setting to lay out the procedure of the online calibration.

With the fixed volatility ratio, we can express the t-dependent factor volatilities in terms of the constant volatilities vector and a t-dependent scaling factor $\alpha(t)$, i.e.

$$p_i(t) = p_i \alpha(t)$$

We assume the scalar function $\alpha(t)$ is a piecewise constant function, with the time knots being the standard option expiration dates t_1, t_2, t_3, \dots . Namely,

$$\alpha(t) = \alpha_k, \text{ for } t_{k-1} \leq t < t_k \text{ (where } t_0 = 0 \text{)}$$

The log variance of a futures contract with maturity T_M up to its standard option expiry t_M reduces to

$$\begin{aligned} \langle \log F^{T_M} \rangle_0^{t_M} &= \sum_{i,j=1}^N \int_0^{t_M} p_i(s) p_j(s) e^{-(\beta_i + \beta_j)(T_M - s)} \rho_{i,j} ds \\ &= \sum_{i,j=1}^N p_i p_j \rho_{i,j} e^{-(\beta_i + \beta_j)T_M} \sum_{k=1}^M \alpha_k^2 \frac{e^{(\beta_i + \beta_j)t_k} - e^{(\beta_i + \beta_j)t_{k-1}}}{\beta_i + \beta_j} \end{aligned}$$

The price of the first option (with the nearest expiration) depends only on α_1 , and so α_1 can be backed out from the option price. The price of the second option depends on both α_1 and α_2 . Since α_1 is already determined, we just need to solve for α_2 to match the second option. We continue this procedure to solve for the values for all α_k .

The difference with the seasonal model is to use T-dependent volatility functions. Analogous to the non-seasonal model, the factor volatility functions can be expressed as a constant vector multiplied by a T-dependent scalar, namely:

$$q_i(T) = q_i \lambda(T)$$

The log variance of a futures contract with maturity T_M up to its standard option expiry t_M is equal to

$$\begin{aligned}
\langle \log F^{T_M} \rangle_0^{t_M} &= \sum_{i,j=1}^N \int_0^{t_M} q_i(T_M) q_j(T_M) e^{-(\beta_i + \beta_j)(T_M - s)} \rho_{i,j} ds \\
&= \lambda^2(T_M) \sum_{i,j=1}^N q_i q_j \rho_{i,j} e^{-(\beta_i + \beta_j)T_M} \frac{e^{(\beta_i + \beta_j)t_M} - 1}{\beta_i + \beta_j}
\end{aligned}$$

This is easier to solve than the non-seasonal model and no bootstrapping procedure is required.

The instruments used in this calibration are vanilla options on futures, i.e., vanilla options expiring in t_M on the futures expiring on T_M . The vanilla option volatilities used in the calibration are marked by the trading desk.

There are two main sources of swaption market prices that we use to gauge how other deals are valuing the swaptions. The first source is the inter-dealer broker market. The second source is the monthly FinPricing swaption valuation service that provides a consensus (average) price among the participants of this service.

In practice, we would like to hold the mean reversion, volatility ratio and correlation parameters as steady as possible. Changing these parameters will change the Greeks, and requires re-balancing of the futures and options hedges, hence incurring costs. Therefore, we will only change these parameters when they noticeably deviate from both FinPricing and broker quotes. More specifically, we consider re-calibration if the model price is more aggressive than both FinPricing and broker offers. We also consider re-calibration if the model price is more conservative than FinPricing and broker bids.

Let me use a real-life example to illustrate this process. At Feb 2021 month end, we have the following comparison between the model price and FinPricing consensus for WTI swaptions. The model prices are within 1% of FinPricing for Cal24, 25, and 26. For cal22 and 23, the model price is conservative by more than 1%. Shall we recalibrate the model to show higher prices for Cal22 and 23?

Straddle Swaption	Model Price	FinPricing Consensus	Diff %
Cal22	9.49	9.97	-4.8%
Cal23	12.86	13.07	-1.6%
Cal24	15.07	15.13	-0.4%
Cal25	16.69	16.71	-0.1%
Cal26	17.87	17.8	0.4%

Table 1: The comparison of the model price and market price for WTI swaptions

In the next few weeks, we got the following broker quotes and compared with our model. Clearly, our model is higher than the broker bids. So, we are fine with the model being conservative relative to FinPricing, knowing that we are not over-conservative (which would be the case if our model is even lower than the broker bid). No re-calibration is needed.

Instrument	Quote Date	Ref Swap	Bid	Offer	Model Price
Cal22 ATM Call	3/19/2021	88.00	4.57	4.97	4.75
Cal23 ATM Call	3/19/2021	83.50	6.05	6.45	6.49
Cal24 ATM Call	3/24/2021	81.75	6.93	7.33	7.42
Cal25 ATM Call	3/24/2021	80.50	7.60	8.10	8.22
Cal26 ATM Call	3/24/2021	80.00	8.26	8.76	8.87

Table 2: The comparison between broker quotes and model prices.

Finally, if the model needs to be recalibrated, we tune the model parameters (mean reversion, vol ratio, and correlation) to put the swaption prices back in the acceptable (as determined by the procedure above) range again. Please refer to the following table for the sensitivity to the model parameters. Tweaking the mean reversion higher has an impact of reducing the swaption price in the front tenors, and increasing the swaption price in the back tenors.

Tweaking the vol ratio higher has an impact of lowering the swaption price across the board. Tweaking the correlation higher reduces the swaption value across all tenors, but the impact is higher for the swaptions in the back. One can use the 3 parameters to achieve different “shape” changes. For example, if one wants to lower the swaption values by similar amount across tenors. The most effective tool is to increase the vol ratio. Mean reversion and correlation provide further flexibility for tuning the “shape”.

ATM Call	Base Price	Mean Rev + 0.1	Diff to Base	Vol Ratio + 0.1	Diff to Base	Corr + 10%	Diff to Base
Cal22	4.746	4.724	(0.022)	4.688	(0.059)	4.730	(0.016)
Cal23	6.430	6.428	(0.002)	6.358	(0.072)	6.389	(0.041)
Cal24	7.533	7.542	0.009	7.460	(0.073)	7.478	(0.055)
Cal25	8.347	8.360	0.013	8.277	(0.070)	8.287	(0.060)
Cal26	8.934	8.949	0.015	8.870	(0.065)	8.873	(0.061)

Table 3: The impact of parameters on prices

4. Pricing Commodity Derivatives

Asian option settles on the average of the futures price over a sampling period. Typically, the prompt futures contract is used as the sampling contract. Since the arithmetic average of the lognormal variables is no longer lognormal, one typically utilizes some approximation to price Asian options. The most common method is to assume that the arithmetic average is still a lognormal variable and to use moment matching to determine the variance of this average.

Suppose the Asian option has M sampling dates t_1, t_2, \dots, t_M , with sampling weights w_1, w_2, \dots, w_M . The average price for which we need to calculate the log variance is the following:

$$F = \sum_{k=1}^M w_k F^{T_k}(t_k)$$

By the moment matching approximation, the log variance of F can be solved as

$$\mathbb{V}[\log F] = \log \frac{\mathbb{E}[F^2]}{\mathbb{E}^2[F]}$$

The expected value in the numerator can be further expanded as follows

$$\begin{aligned} \mathbb{E}[F^2] &= \sum_{j,k=1}^M w_j w_k \mathbb{E}[F^{T_j}(t_j) F^{T_k}(t_k)] \\ &= \sum_{j,k=1}^M w_j w_k \mathbb{E}[F^{T_j}(t_j)] \mathbb{E}[F^{T_k}(t_k)] e^{\langle \log F^{T_j}, \log F^{T_k} \rangle_0^{\text{Min}(t_j, t_k)}} \end{aligned}$$

Finally, the log variance is plugged into the Black-Scholes equation to arrive at the Asian option price.

Another example is commodity swaption. A swaption is like Asian option as it is also an option on the average of futures prices. The difference is that a swaption expires before the averaging period starts. The formula to price a swaption using the multi-factor model is very similar to the Asian option formula above, except that the sampling dates t_1, t_2, \dots, t_M are all the same and equal to the swaption expiration date.

Moreover, for OTC swaptions, the discount factors from the swap settlement dates to the swaption expiration date are baked into the sampling weights. More specifically, if the undiscounted sampling weights for the underlying futures are (u_1, u_2, \dots, u_M) , and the discount factors from the swaption expiration date to the settlement dates of each underlying swaplet are (d_1, d_2, \dots, d_M) , then the discounted sampling weights which are used in swaption pricing are

$$(w_1, w_2, \dots, w_M) = (u_1 * d_1, u_2 * d_2, \dots, u_M * d_M)$$

So far, we have covered the pricing formula for a native denominated (USD typically) Asian option or Swaption. We now discuss how to price Asian or Swaption denominated in a foreign currency. To this end, we first review the stochastic calculus associated with a foreign currency.

Let's denote by d the domestic economy (typically USD), and by f the foreign economy. Let's use P to stand for the domestic risk-free measure, and \tilde{P} for the foreign measure. $X(t)$ is the spot FX rate at time t to convert 1 unit of foreign currency into the domestic currency. $Y(t) = 1/X(t)$ is the spot FX rate that converts 1 unit of the domestic currency to the foreign currency. As usual, $F^T(t)$ denotes the futures price in the domestic currency, and we define $G^T(t) = F^T(t)Y(t)$ to be the futures price in the foreign currency.

The dynamics of $X(t)$ is governed by the following SDE, where the Brownian motion W is in the domestic measure P ,

$$\frac{dX(t)}{X(t)} = (r_d(t) - r_f(t)) dt + \sigma_X(t) dW(t)$$

Here, $r_d(t)$ and $r_f(t)$ are the deterministic interest rates for the domestic and foreign currencies respectively. By Ito lemma,

$$dY(t) = -\frac{dX}{X^2} + \frac{\langle dX \rangle}{X^3} = Y(t) \left[(r_f(t) - r_d(t)) dt - \sigma_X(t) (dW(t) - \sigma_X(t) dt) \right]$$

By the Girsanov theorem, we can rewrite this SDE in the foreign measure \tilde{P} , where $d\tilde{W}(t) = dW(t) - \sigma_X(t) dt$ is a Brownian motion in the foreign measure,

$$\frac{dY(t)}{Y(t)} = (r_f(t) - r_d(t)) dt - \sigma_X(t) d\tilde{W}(t)$$

Therefore, Y and X have the same volatility and covariance.

Consider an Asian option that settles on the average price in the foreign currency:

$$G = \sum_{k=1}^M w_k G^{T_k}(t_k) = \sum_{k=1}^M w_k F^{T_k}(t_k) Y(t_k)$$

It is convenient to price the foreign option in the foreign measure \tilde{P} , and we use the same moment matching method as the native option,

$$\mathbb{V}[\log G] = \log \frac{\mathbb{E}_{\tilde{P}}[G^2]}{\mathbb{E}_{\tilde{P}}^2[G]}$$

The first moment is easy to compute

$$\mathbb{E}_{\tilde{P}}[G] = \sum_{k=1}^M w_k \mathbb{E}_P[F^{T_k}(t_k)] \mathbb{E}_{\tilde{P}}[Y(t_k)]$$

The first expectation in the right-hand side of the equation is simply today's futures price, and the second expectation is today's forward FX rate.

The second moment can be expanded as

$$\begin{aligned} \mathbb{E}_{\tilde{P}}[G^2] &= \sum_{j,k=1}^M w_j w_k \mathbb{E}_{\tilde{P}}[G^{T_j}(t_j) G^{T_k}(t_k)] \\ &= \sum_{j,k=1}^M w_j w_k \mathbb{E}_{\tilde{P}}[G^{T_j}(t_j)] \mathbb{E}_{\tilde{P}}[G^{T_k}(t_k)] e^{\langle \log G^{T_j}, \log G^{T_k} \rangle_0^{\text{Min}(t_j, t_k)}} \end{aligned}$$

The expected values are computed the same way as in the first moment. What remains is to calculate the log-covariance term:

$$\begin{aligned} &\langle \log G^{T_j}, \log G^{T_k} \rangle_0^{\text{Min}(t_j, t_k)} \\ &= \langle \log F^{T_j}, \log F^{T_k} \rangle_0^{\text{Min}(t_j, t_k)} + \langle \log F^{T_j}, \log Y \rangle_0^{\text{Min}(t_j, t_k)} \\ &\quad + \langle \log Y, \log F^{T_k} \rangle_0^{\text{Min}(t_j, t_k)} + \langle \log Y, \log Y \rangle_0^{\text{Min}(t_j, t_k)} \end{aligned}$$

Since the FX process is modeled as a single-factor model, the log-covariance terms in the right-hand equation are a particular case of the cross-asset log-covariance formula.

The formula to price a foreign denominated swaption using the multi-factor model is very similar to the Asian option formula above, except that the sampling dates t_1, t_2, \dots, t_M are all the same and equal to the swaption expiration date. Moreover, for OTC swaptions, the discount factors (for the foreign currency since we are pricing in the foreign measure) from the swap settlement dates to the swaption expiration date are baked into the sampling weights (w_1, w_2, \dots, w_M) .

As described in the model, the task of simulating the entire forward curve boils down to that of simulating N mean reverting factors Y . In this section, we describe the details of the numerical algorithm to generate paths for these N mean reverting factors. With these mean reverting factors, the futures prices can be recovered easily using the algebraic formula above.

Suppose we want to simulate the mean reverting factors on M sampling dates t_1, t_2, \dots, t_M , and we have already generated the paths up to time t_{k-1} . To simulate the next time step, we note that

$$Y_i(t_k) = e^{-\beta_i(t_k - t_{k-1})} Y_i(t_{k-1}) + \int_{t_{k-1}}^{t_k} e^{-\beta_i(t_k - s)} p_i(s) dW_i(s)$$

The N -by- N covariance matrix of the integral term is easily derivable from the Ito isometry:

$$\begin{aligned} & \left\langle \int_{t_{k-1}}^{t_k} e^{-\beta_i(t_k - s)} p_i(s) dW_i(s), \int_{t_{k-1}}^{t_k} e^{-\beta_j(t_k - s)} p_j(s) dW_j(s) \right\rangle \\ &= \int_{t_{k-1}}^{t_k} e^{-(\beta_i + \beta_j)(t_k - s)} p_i(s) p_j(s) \rho_{i,j} ds \end{aligned}$$

We then calculate the Cholesky decomposition of the covariance matrix. Equipped with a random number generator capable of drawing N independent Gaussian random numbers, we can simulate the integral term and hence generate the paths for the mean reverting factors at the next time step t_k .

We may follow the recipe above in a multi-asset simulation. It works theoretically, however in practice this requires calculating the full covariance matrix and its Cholesky decomposition across all assets and all factors at every time step.

To make the multi-asset algorithm complete, we need to discuss how to determine the pairwise correlation $\rho_{i,j}^{a,b}$ between the i 'th factor of the asset and the j 'th factor of the b asset (where a could be equal to b , in which case the correlation is between different factors of the same asset). Obviously, one cannot hope to assign arbitrary values to these correlations and still get a semi-positive-definite (SPD) matrix that can be Cholesky decomposed. The simplest way to ensure SPD is to compute the correlation matrix directly from historical data.

But where do we get historical data on the factors? The factors are a model construct and are not directly observable from the historical data. What we observe is the historical futures prices, and we need to somehow “transform” them into factors. Here's a simple way to do this.

Take the two-factor model for WTI as an example. We pick two nearby contracts, one from the front end of the curve, say the 3rd nearby, i.e., $T_1 = t + 3/12$, and other from the back end of the curve, say the 36th nearby, i.e., $T_2 = t + 36/12$. According to the model,

$$\begin{aligned}\frac{dF^{T_1}(t)}{F^{T_1}(t)} &= \sigma(t)[p_S e^{-\beta(T_1-t)} dW_S(t) + p_L dW_L(t)] \\ \frac{dF^{T_2}(t)}{F^{T_2}(t)} &= \sigma(t)[p_S e^{-\beta(T_2-t)} dW_S(t) + p_L dW_L(t)]\end{aligned}$$

The left-hand sides of the equations, the log returns of the two futures, are directly observable from the historical time series. We can solve for the factors as follows:

$$\begin{aligned}p_S \sigma(t) dW_S(t) &= \frac{\frac{dF^{T_1}(t)}{F^{T_1}(t)} - \frac{dF^{T_2}(t)}{F^{T_2}(t)}}{e^{-\beta(T_1-t)} - e^{-\beta(T_2-t)}} \\ p_L \sigma(t) dW_L(t) &= \frac{e^{-\beta(T_2-T_1)} \frac{dF^{T_1}(t)}{F^{T_1}(t)} - \frac{dF^{T_2}(t)}{F^{T_2}(t)}}{e^{-\beta(T_2-T_1)} - 1}\end{aligned}$$

Note that in this formula, the mean reversion rate is already determined by the procedure in Model Calibration. From this, the historical correlation between any two factors (either from the same asset or different assets) can be calculated.

5. Empirical and Numerical Study

Let's go through an example to understand the consequence of choosing a calibration strategy. Let the valuation date be October 2020. Suppose we need to price an option on the December 2022 futures contract, while the option expires on the option expiration of the December 2021 contract. This is what we call an Early Expiry Option, which is a mini-swaption. T

The key element in pricing this option is the early expiry volatility of the December 2022 contract during the period from now to the expiration of the December 2021 contract. If we use a seasonal model to price this option, the early volatility of the option depends entirely on the trader's volatility mark for the December 2022 contract. If we use a non-seasonal model to price this option, the early volatility depends entirely on the trader's volatility mark for December 2021 contract. In other words, the Vega of the option would be attributed to different contract months depending on which calibration strategy to use.

Now let's imagine a market scenario where the implied volatility of the December 2021 contract increased by 2%, while the volatility of the December 2022 stays put. If we use the seasonal model, the Vega profit and loss (PnL) under this market scenario would be flat. But if we use the non-seasonal model, the Vega PnL under this market scenario would be positive, if we are long this option.

The intuition is that the seasonal model interprets this scenario as saying that even though the December 2021 contract becomes more volatile, it doesn't say anything about the December 2022 contract. On the other hand, the non-seasonal model interprets the market scenario as saying that not only December 2021 contract becomes more volatile, but December 2022 contract also becomes more volatile, during the period from now to the expiration of the 2021 contract.

How can the implied volatility of the December 2022 contract not change? It must be that the volatility of that contract during the period from the expiration of 2021 contract to its own expiration

has decreased, to balance the increase of the volatility in its earlier period so that the average volatility for the whole lifetime is still the same.

Which interpretation is correct? Both satisfy the requirement of matching the model to the market implied volatilities for vanilla options. But both are only approximations to what's happening in the real world, and neither is absolutely correct under all circumstances. The truth often lies in between. So, we're interested in seeking a hybrid model that combines the seasonal and non-seasonal calibration strategy. Equipped with this extra flexibility, we will be more capable of capturing the responsive behavior of the market swaption prices, when the term structure of implied volatilities changes.

Below are details on how to calibrate a hybrid model. Let's assume that the volatility ratio is given by the constant vector $p_i = q_i$, and we use the symbol ε to denote the degree to which we want the hybrid model to be more like the seasonal or non-seasonal model. This “non-fungibility” parameter is a number between 0 and 1, and the hybrid model reduces to the seasonal model when $\varepsilon = 1$ and it becomes the non-seasonal model when $\varepsilon = 0$. Let's also denote the implied volatility, the option expiration and futures expiration of the M'th contract as σ_M, t_M, T_M respectively.

The first step is to calibrate a fully seasonal model, and then assign the result “partially” to the T-dependent local volatility function $q(T)$. For this purpose, we recall that in a seasonal model, the scalar λ satisfies the equation:

$$\begin{aligned}\sigma_M^2 t_M &= \sum_{i,j=1}^N \int_0^{t_M} q_i(T_M) q_j(T_M) e^{-(\beta_i + \beta_j)(T_M - s)} \rho_{i,j} ds \\ &= \lambda^2(T_M) \sum_{i,j=1}^N q_i q_j \rho_{i,j} e^{-(\beta_i + \beta_j)T_M} \frac{e^{(\beta_i + \beta_j)t_M} - 1}{\beta_i + \beta_j}\end{aligned}$$

So

$$\lambda(T_M) = \sigma_M \sqrt{\frac{t_M}{\sum_{i,j=1}^N q_i q_j \rho_{i,j} e^{-(\beta_i + \beta_j) T_M} \frac{e^{(\beta_i + \beta_j) t_M} - 1}{\beta_i + \beta_j}}}$$

Now we set the T-dependent local volatility to be a fraction (in the log space) of the solution above, according to the non-fungibility parameter:

$$q_i(T) = q_i^\varepsilon \lambda^\varepsilon(T)$$

The “residual” fungible implied volatilities $\sigma_M / \lambda^\varepsilon(T_M)$ are then to be matched using a boot-strapping approach., the same way as in the non-seasonal model. Recall that the equations to solve for are

$$\begin{aligned} \sigma_M^2 t_M / \lambda^{2\varepsilon}(T_M) &= \sum_{i,j=1}^N \int_0^{t_M} p_i(s) p_j(s) e^{-(\beta_i + \beta_j)(T_M - s)} \rho_{i,j} ds \\ &= \sum_{i,j=1}^N p_i p_j \rho_{i,j} e^{-(\beta_i + \beta_j) T_M} \sum_{k=1}^M \alpha_k^2 \frac{e^{(\beta_i + \beta_j) t_k} - e^{(\beta_i + \beta_j) t_{k-1}}}{\beta_i + \beta_j} \end{aligned}$$

We set the t-dependent local volatility to be

$$p_i(t) = p_i^{1-\varepsilon} \alpha(T)$$

It is easy to verify that such calibrated model always matches the implied volatilities of all vanilla options, i.e.

$$\sigma_M^2 t_M = \sum_{i,j=1}^N \int_0^{t_M} p_i(s) p_j(s) q_i(T_M) q_j(T_M) e^{-(\beta_i + \beta_j)(T_M - s)} \rho_{i,j} ds$$

When $\varepsilon = 0$, the model is equivalent to the fully non-seasonal model, and when $\varepsilon = 1$, the model is equivalent to the fully seasonal model. For the early expiration option example earlier, when $\varepsilon = 1/2$, we will have about half of the vega attributed to December 2021 contract and half to the December 2022 contract.

Finally, how do we calibrate the non-fungibility parameter ε in the model? Since the main impact of the parameter is to the Vega attribution, we can monitor market swaption prices over a period of time where the term structure of implied volatilities goes through significant changes and choose ε such that the responsive behavior of the model matches the market prices the best. This can be tricky though since there could be other factors affecting the changes in swaption prices, especially over an extended period.

As an illustration, we pick four month-end dates in 2020, and calculate the IPV (difference between model price and FinPricing consensus) for our existing WTI swaption positions, using 3 different models: fully non-seasonal, fully seasonal and hybrid with 50% non-fungibility parameter. The model parameters are chosen such that the swaption NPV as of October 8, 20 are about the same for the 3 models. One can see that for this 4-month period, the hybrid model is more stable than the non-seasonal model, but it is hard to judge between the seasonal and hybrid models. But for oil assets, the hybrid model is naturally preferred over the fully seasonal model.




Figure 5. Month-end term structure of implied volatilities.

Model Params	Non-Seasonal	Seasonal	Hybrid
Mean Reversion	0.59	0.53	0.56
Volatility Ratio	1.32	1.67	1.44
Correlation	0%	0%	0%
Non-Fungibility	0%	100%	50%

Table 4. Calibrated model parameters for seasonal, non-seasonal, or hybrid assumptions

WTI Swaption IPV	Non-Seasonal	Seasonal	Hybrid
6/28/2020	\$ 432,064	\$ 504,492	\$ 458,317
7/31/2020	\$ 444,942	\$ 695,192	\$ 561,262
8/30/2020	\$ (25,346)	\$ 635,440	\$ 300,675
9/30/2020	\$ 346,722	\$ 555,789	\$ 442,195

Table 5. Valuations results under seasonal, non-seasonal or hybrid assumptions




Figure 6. The plots of month-end swaption IPVs

The “curve” risk associated with only the swaptions are exotic in nature and cannot be hedged with vanilla options. The “curve” risk means the uncertainty with how the front and back end of the futures curve tend to move together. If the front end and back end moves in tandem, swaptions should increase in value relative to vanilla options. If the front and the back move in opposite directions, swaptions should decrease in value relative to vanilla options.

There’s no way to offset this curve risk effectively with vanilla options. Of course, the most effective way to hedge this is to sell the swaptions we bought. But it’s quite often that we cannot do this without incurring big costs. Another less effective way is to use time spread options as a hedge. But time spread options are as illiquid as swaptions.

The manifestation of this in the model is that the risks to the model parameters, including volatility ratio, mean reversion and correlations are, to the most part, un-hedgable (unless one sells the swaptions he bought). These parameters are not directly observable from the market either. So how to mark these parameters and manage the uncertainty with these parameters presents a challenge to the model. But there are at least 3 types of mitigation to this challenge.

First, we impose a hard limit on how much swaption Vega traders can warehouse. This limits our exposure to the un-hedgable swaption risks.

Second, we verify the model valuation at least monthly. This ensures that we mark the model to the market expectation of how the “curve” dynamics will actualize.

Finally, we review the historical analysis periodically to ensure we are not stepping into an overheated market. But typically mark-to-market is the rule rather than exception. No amount of historical analysis can tell us for sure what is going to happen in the future.

6. Conclusions

The article presents a model for pricing commodity derivatives, such as commodity swaptions. Swaptions are a relatively illiquid product in commodities market, and the natural flow tends to be one-sided. As a result, dealers typically warehouse long swaption risks. Delta and Vega of these swaptions can be effectively hedged with liquid vanilla options and futures.

We divide the model calibration strategy into seasonal and non-seasonal. Intuitively, the futures contracts for a seasonal asset, e.g., power or gas, are less fungible than futures contracts for a non-seasonal asset, e.g., oil. Hence, we use a bootstrapping strategy to calibrate local volatilities for non-seasonal assets, but for seasonal assets we calibrate the local volatilities of each contract separately.

The model depending on multi factors can capture price dynamics and volatility more precisely, that is of great importance to commodity market participants in order to exercise sound trading, hedging, and risk-managing strategies.

We decompose the historical daily returns of a futures curve into principal components and compare the first two principal components with those implied by the two-factor model. The mean reversion

rate, volatility ratio and correlation are chosen to achieve the best fit to the historical principal components in the least squared sense.

Empirically we find the model prices match historical market prices very well, indicating the model is quite accurate. We also discuss the model risk as well as the curve risk, and propose some mitigation solutions.

References:

Bakshi, G. et al., “Understanding the sources of risk underlying the cross section of commodity returns,” *Management Science* (2019)

Benth, F. and Saltyte-Benth, J., “The normal inverse Gaussian distribution and spot price modelling in energy markets.” *International Journal of Theoretical Applied Finance*, 7(02) (2004), pp. 177-192

Borovkova, S., Bunk, H., Goeij, W., Mechev, D. and Veldhuizen, D., “Collagenized commodity obligations.” *Journal of Alternative Investments*, 16 (2) (2013), pp. 9-29

Borovkova, S. and Geman, H., “Seasonal and stochastic effects in commodity forward curves.” *Review of Derivatives Research*, 9 (2) (2006), pp. 167-186

Cortazar, G. et al., “Time-varying term structure of oil risk premia,” *The Energy Journal*, 43 (5) (2022), 71-91.

Cristoffersen, P., Lunde A., and Olesen, K., “Factor structure in commodity futures return and volatility,” *Journal of Financial and Quantitative Analysis*, 54 (3) (2019), 1083-1115.

Fasanya, I. and Awodimila, C., “Are commodity prices good predictors of inflation? The African perspective,” *Resources Policy*, 69 (2020): 101802.

George, C., :Estimating the term structure of commodity market preferences,” *European Journal of Operational Research*, 282 (3) (2020), 1146-1163.

Hollstein, F., Prokopczuk, M., and Wuersig, C., “Volatility term structures in commodity markets,” *Journal of Futures Markets*, 40 (4) (2020), 527-555.

Hu, C., Li, Z., and Liu X., “Liquidity shocks, commodity financialization, and market co-movements,” *Journal of Futures Markets*, 40 (2020): 1315-36.

Idilbi-Bayaa, Y. and Qadan, M., “Forecasting commodity prices using the term structure,” *Journal of Risk and Financial Management*, 14 (2021): 585.

Kirk, E., “Correlation in the energy markets. Managing Energy Price Risk.” *Risk Publications and Enron* (1995), pp. 71-78

Ladokhin, S. and Borovkova, S. “Three-factor commodity forward curve model and its joint P and Q dynamics.” *Energy Economics*, 101 (2021) 105-418

Schwartz, E., & Smith, J. E. (2000). Short-Term Variations and Long-Term Dynamics in Commodity Prices. *Management Science*, 46(7), 893-911.

Schwartz E. “The stochastic behavior of commodity prices: Implications for valuation and hedging.”
J. Finance, 52 (3) (1997), pp. 923-973

Trolle, A. and Schwartz, E. “Unpanned stochastic volatility and the pricing of commodity derivatives.” *Review of Financial Study*, 22 (2009), pp. 4423-4461